Beyond the Cost Model: Understanding Price Elasticity and Its Applications

Serhat Guven, FCAS, MAAA, and Michael McPhail, FCAS, MAAA

Abstract

Once cost models have been constructed, insurers spend a significant amount of time translating those expected cost models into a rating algorithm. Today, competitive analytics are widely used to support this effort. However, companies often fail to fully integrate competitive analytics into the pricing process. The intent of this paper is to provide the basic tools needed for insurers to make more effective pricing decisions using customer price elasticity of demand. To achieve this, we will explore demand modeling techniques, as well as practical applications of demand modeling in pricing.

Keywords. Price elasticity; demand; generalized linear modeling; price simulation; price optimization

1. INTRODUCTION

We learned in our introductory economics courses that price elasticity of demand (PED) is loosely defined as the change in demand for a given change in price. It measures a consumer's sensitivity to a change in price for a given good or service, ranging from high sensitivity (elastic) to low sensitivity (inelastic). In equation form,

$$E = -\frac{(\mu_{P1}-\mu_{P0})/\mu_{P0}}{(P1-P0)/P0} \tag{1.1}$$

Where E = PED

$P0$ = Initial Price

$P1$ = New Price

μ_{P0} = Demand at initial price

μ_{P1} = Demand at new price

One of the benefits of PED is that it enables a firm to enhance its pricing strategies through a better understanding of the price sensitivity of its target market. Let's look at a basic example. Suppose an accountant has been operating a private firm for the past five years and has built a base of loyal customers. Although sales are stable, growth is stagnant and he is becoming concerned with the long-term viability of the firm. He decides to implement a 20% discount on first-time tax returns to bring in new business. As an accountant, he understands the need to maintain profitability and simultaneously raises the price of a tax return from \$250 to \$275. Knowing that his current customers are extremely loyal, he believes the majority of them will accept the slight increase rather than face the risks associated with switching accountants (i.e. lower quality). In other words, he is altering his pricing strategy based on the price elasticity of his target market. For new business, which tends to be more price sensitive, he is lowering the price of first-time tax returns to $220 ($275 x 80%). He offsets the discount by raising the price for his existing business, which tends to be less price sensitive. Through this pricing strategy, he expects to increase his customer base and revenue in the long term. In reality, there are a number of factors that must be considered when altering pricing strategy, but hopefully this example provides some insight into how PED can be utilized.

Another benefit of PED is the accurate measurement of revenue changes in response to pricing changes. In the above example, the accountant increased the price of tax returns for current customers by 10% (275/250 - 1). Although a majority of customers will accept the increase, some customers are expected to switch accountants. Ignoring growth, consider the following:

(1) Current Price	(2) Proposed Price	(3) # of Current Customers	(4) # of Customers Expected to Leave Due to Rate Change	(5) = (1) * (3) Expected Revenue Under Current Rates	(6) = (2) * ((3) – (4)) Expected Revenue Under Proposed Rates	(7) = (6)/(5) - 1 Expected % Change in Revenue
$250	$275	100	5	$25,000	$26,125	4.5%

Table 1.1 Price change vs. revenue change

Although the accountant is still profitable, the expected percentage change in revenue is less than half of the proposed rate change. But what if the number of customers expected to leave were ten rather than five? In that case, the expected percentage change in revenue would actually be negative, resulting in an overall loss to the business. Understanding the effects of demand on revenue is crucial to making pricing decisions that improve the overall position of the firm.

As an actuary in the P&C insurance market, PED opens up a whole new realm of possibilities. Although competitive analytics are widely used in the industry today, many companies fail to fully integrate them into the pricing process. For a specific firm, what does it mean to be competitive? Is it purely a function of price or should more abstract concepts such as brand and loyalty also be considered? How can a firm systematically incorporate that information into its pricing strategy? The rest of this paper will attempt to answer these questions, with a focus on basic techniques that can be used to model demand. In Section 2, we discuss demand modeling including model form and model structure. In Sections 3 & 4, we explore practical applications of PED in the context of price simulation and price optimization. Finally, in Section 5, we provide our concluding remarks.

Before moving into demand modeling, the authors would like to provide a clarification on the intent of this paper. Due to minimal actuarial literature regarding PED, this paper strives to offer an introductory review of the topic as opposed to an actual practitioner's guide. As such, concepts deserving of numerous pages receive little treatment for the sake of brevity. Our hope is that this paper provides a springboard for further work on PED.

2. DEMAND MODELING

In statistical models, there are two basic types of data variables. The fitted response variable is the output of the model whereas predictor variables are inputs to the model. Given a set of predictors, the model will provide a fitted response. The actual response in the data has two elements. The signal represents the predictive behavior and the noise is the random behavior. The goal of the model is to separate the signal from the noise. Thus, we are trying to identify a function that, when applied to the predictors, produces a fitted response which represents the signal in the actual response.

When modeling the price elasticity of demand, the first task is to clearly define the actual response. In addition, it is useful to classify the different types of predictors that are to be studied. There are significant challenges in clarifying the data for elasticity modeling. For example, many companies wrestle with the definition of a quote – questions such as how long a quote can stay open before it is considered a rejection, how do you deal with multiple changes to endorsements, among many other such issues. These issues are outside of the scope of this paper.

Price elasticity modeling builds a model in which the actual response is the individual customer's acceptance or rejection of a quote or renewal offering. The predictors in these models can be classified as price and non price related. The fitted elasticity is derived from the coefficients that are associated with the price related predictors in the model. Note it can be useful to further categorize price factors as external (reflecting pricing activities of competitors) and internal (reflecting pricing activities of the company).

In the following sections, we will be discussing the structural form of a model where, given the actual response and predictors, we will be able to calculate fitted responses and derive fitted elasticity. In discussing the model form, we will focus on the following three areas:

- Error distribution of the actual response
- Functional form that relates the actual response to the predictors
- Structural design of the predictors

Furthermore, we will discuss key validation techniques to assess the quality of the projections. The result of the form and modeling exercise is to improve understanding of a consumer's behavior toward price.

2.1. Error Distribution

Unlike many other responses that practitioners model, it is a relatively straight forward exercise to select the error distribution function of a price elasticity model. Since the actual response is of the form yes/no, the binomial distribution is the most appropriate choice. Note that it is possible to have intermediate phases that are associated with different degrees of acceptance or rejection, but that is outside the scope of this article.

Before continuing, it is useful to understand what it means to have a binomial distribution. Error distributions are defined in empirical terms as well as functional forms. Empirically, the actual response can only be 0 or 1, meaning the selection of a binomial structure is clearly supported. The functional form of an error distribution relates the variance of the actual response to the fitted response. Another interesting interpretation is to understand how the variance function plays a part in model form. The relationship of the variance to the mean in a binomial distribution assumes the following form:

Variance = Mean x (1 – Mean)

Because a binary response is modeled, the fitted response will lie between 0 and 1 (more on this later). Given that constraint, the following picture shows the relationship between the mean and the variance:

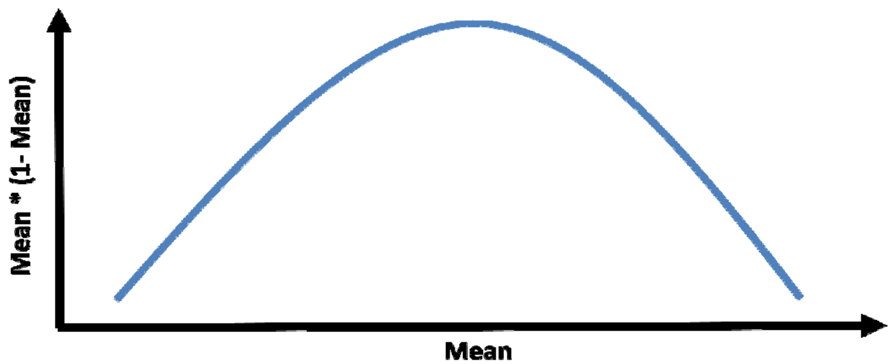

Figure 2.1 Variance mean relationship of binomial distribution

The variance of a result represents the degree of uncertainty associated with the predictions of that result. In a cost model, it is common to assume that the larger the fitted value, the greater the uncertainty (variability) of the estimate. In elasticity models, both large and small estimates have lower variability whereas mid-range estimates have higher variability. This is the inherent nature of binomial models. If all the records are always rejecting offers (actual response is low - zero), then the modeler can feel fairly confident that future offers will be rejected. Similarly, if all records are always accepting offers, (actual response is high - one) then the same high level of confidence can be said about forecasted results. Now if the actual response is mid range, i.e. 50%, then there would be much greater uncertainty associated with the fitted value. Once the distribution function has been selected, the next key structure is the link function.

2.2. Link Function

The link function is the functional form that relates the response to the predictors. There are two key requirements when selecting a link function for elasticity modeling. The first is that the fitted response of the probability of acceptance or rejection of a price offer should lie between 0 and 1. The requirement of the link function is to transform the resulting coefficients from the structural design into a result that is consistent with the laws of probability. Note that in common statistical terminology, the combination of coefficients from the structural design of the model is called a linear predictor. It is not necessary for the linear predictor to be purely linear, as will be discussed in following sections.

The second requirement, though less stringent, is that the fitted responses approach pure acceptance or rejection but never actually reach it. This asymptotic behavior is best visualized as the following S-shape curve:

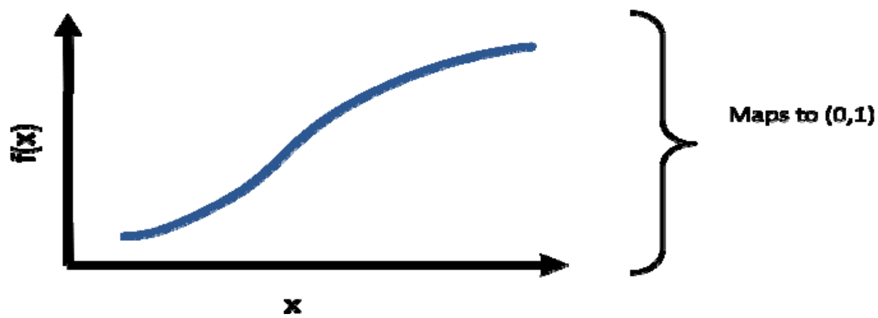

Figure 2.2 Shape of modeled response for elasticity modeling

Allowing the fitted response to reach pure acceptance or rejection implies too much certainty in the forecast.

Unlike error distributions, there are many choices of link functions that satisfy these requirements. Two functions commonly used in practice are:

1. The logit or logistic function

$$\mu = f(x) = \frac{1}{1+1/\exp(x)} \tag{2.1}$$

   Where x is the linear predictor

2. The probit or normal function

$$\mu = f(x) = \Phi(x) \tag{2.2}$$

   Where Φ is the cumulative normal distribution

When selecting between different functions, it can be useful to perform validation tests using the elasticity concept. Recall that because we cannot observe individual elasticity, we have to derive it from the fitted responses from the model. There are many ways to define elasticity, and we will focus on the following two:

1. Classical Elasticity, which is the percentage change in demand due to the percentage change in price. As defined in the introduction, for a given customer,

$$E = -\frac{(\mu_{P1} - \mu_{P0})/\mu_{P0}}{(P1 - P0)/P0} \tag{2.3}$$

   Assume P1 > P0 so that the numerator is negative and the denominator is positive. From this formula, the initial expected demand is 'fixed' so elasticity has a linear relationship with demand. Note that μ_{P_i} is the expected probability of success associated with the price P_i.

2. Linear Predictor Elasticity, which we define as the change in the linear predictor due to the percentage change in price. Specifically, for a given customer:

$$E = -\frac{\left(\beta_0 + \propto \times \frac{P1}{P0}\right) - \left(\beta_0 + \propto \times \frac{P0}{P0}\right)}{(P1 - P0)/P0} \tag{2.4}$$

   Assuming a simple linear predictor with no nonlinear components and one price factor:

$$E = -\propto \tag{2.5}$$

   Where α is the coefficient associated with the price factor

   Note that in this simple case, E does not vary with demand. Thus, as the expected demand increases, we expect E to be constant. This relationship holds true regardless of the number of non-price factors in the model structure.

Comparing these two definitions with respect to demand yields the following two lines on this theoretical graph:

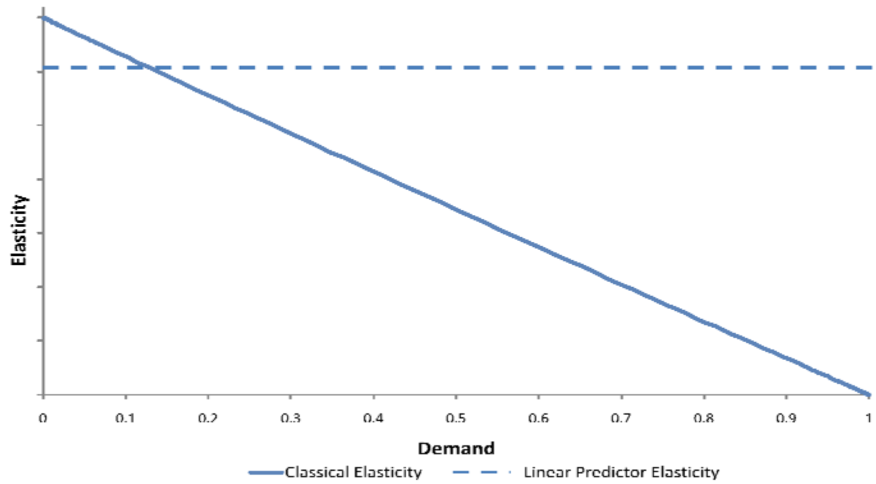

Figure 2.3 Classical and linear predictor elasticity vs. demand

As shown in Figure 2.3, the linear predictor elasticity (dashed) does not vary by demand whereas classical elasticity (solid) has a linear relationship with demand. While there is limited interpretative appeal of linear predictor elasticity, we developed this particular definition because it has useful properties when assessing link functions.

Consider the following validation exercise. A model was built using quote data from months one through six. During months seven and eight, the number of quotes and converted policies were captured. Also assume that a rate change took effect between months seven and eight. The model was used to score the expected close rates of the month seven and month eight quotes. Using these scores, the two months of data were ranked from highest to lowest and placed into 17 bins each having approximately the same number of quotes. For each bin, the average linear predictor elasticity was calculated as follows:

Month 7 Linear Predictor = $\ln(CR_7/(1-CR_7))$

Month 8 Linear Predictor = $\ln(CR_8/(1-CR_8))$

Where CR_i is the Close Rate in month i

Then, the change in the observed linear predictor divided by the change in average price between the two months was calculated. This yields the following graph where the model used a logit link function:

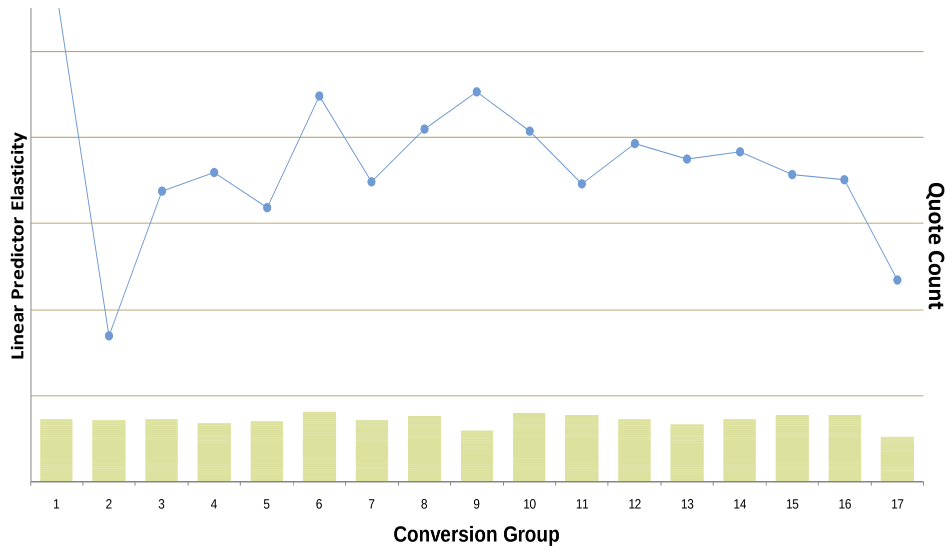

Figure 2.4 Linear predictor elasticity vs. demand for logit link function

In this graph, the x-axis is the expected conversion for the quotes in both months. This score was used to create approximately equal weighted buckets as identified by the bars from the right hand y-axis. On the left hand y-axis, the average linear predictor elasticity was calculated for each group as the change in the aggregate observed linear predictor divided by the change in average price between months seven and eight.

Building a similar graph with the same model structure using a probit link function produces the following:

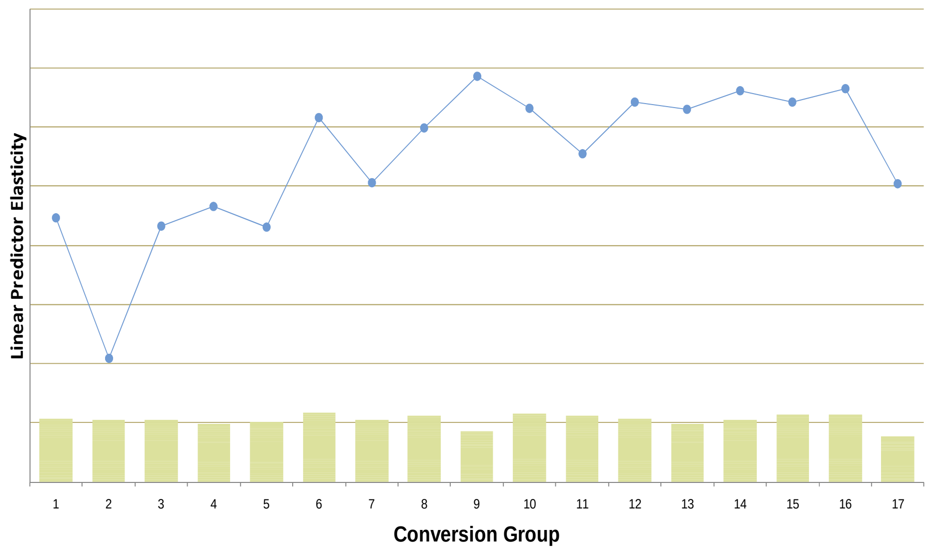

Figure 2.5 Linear predictor elasticity vs. demand for probit link function

For this particular data set, the actual linear predictor elasticity line in the logistic model is flatter than the actual linear predictor elasticity line in the probit model. From the prior theoretical discussion, the linear predictor elasticity would be expected to be flat when plotted against the expected demand. In this case, the logistic link function outperforms the probit link function as it results in a linear predictor elasticity that is more constant relative to demand.

As with any modeling exercise, the practitioner's choice should be validated to ensure the best possible model. Now that the distribution and link function have been selected, the next step is to identify the structural design of the predictors.

2.3. Model Structure

Building the model structure is a balance between over-fitting and under-fitting. When the model structure is too complex, there is a greater likelihood of over-fitting the data. This generally results in the loss of predictive power. However, when the structure is too simple, there is a likelihood of under-fitting the data, also resulting in the loss of explanatory power. When building the model structure, there are a variety of tests the analyst performs to ensure that each element in the model reflects this balance. These include:

- Statistical Tests - these include but are not limited to basic Type III tests, standard errors, etc.
- Consistency Tests – typically these test consistency across different time frames in the data. Alternatively, consistency can be tested across different random data splits.
- Judgment – the best models reflect the balance of technical information and knowledge of the underlying process.

The structural design of the model is often called the linear predictor. It reflects the combinations of parameters or coefficients that are applied to a particular observation. Note that the modeler specifies the structural design, the link function, and error distribution. The model then calculates a set of parameters that are based on these assumptions.

The linear predictor contains a variety of different objects to reflect the different predictors. In addition to categorizing predictors as price and non-price factors, they can also be classified as categorical and continuous elements. When a predictor is considered categorical, a parameter is calculated for each level in the predictor. When a predictor is considered continuous, a parameter is based on the form of the function introduced. The following example illustrates this concept. Without loss of generality, assume the following simple model structure:

$$\text{Base} + \text{RatingArea} + \text{NamedInsuredAge} + f_1(\text{RateChange}) + f_2(\text{CompetitivePosition})$$

In this case, the CompetitivePosition is defined as the MarketPrice/CompanyPrice, where the MarketPrice is a straight average of five competitors. This particular metric was useful in this model. There are many other competitive metrics that can be used in demand modeling (rank, cheapest ratios, etc.) and there are many challenges in obtaining this external data. These concepts are outside of the scope of this paper.

Looking back at the model form in the example, RatingArea and NamedInsuredAge are treated as categorical elements whereas RateChange and CompetitivePosition are continuous elements which are introduced as first degree polynomials. Consider the following output from a model:

Level	Predictor	Component
Base		-0.37876
NamedInsuredAge	25	0.39105
RatingArea	1A	0.07411
RateChange	1.09	-0.9235 * 1.09 + 1.0389
CompetitivePosition	1.07	-0.9414 * 1.07 + 1.1767
Linear Predictor		0.28809
Fitted Value		0.57153

Table 2.1 Model structure example; no interactions

The component represents the parameter applied to the predictor. For categorical factors, the parameter and component are the same. For continuous factors, the parameter is a coefficient applied to the predictor to yield the component. So the linear predictor for this particular observation is the sum of the components. These components are either specific parameters for categorical concepts OR functions of continuous concepts. The fitted value is then the inverse link function applied to the linear predictor. In the example, the model is using a logit link function as defined earlier.

The prior example reflects a model structure in which the predictors are all treated as main effects. This is a strictly linear concept in which the relationship between predictors does not vary based on other predictors. Adding main effects to the model should be done with great care to ensure the main effect is both statistically valid and the resulting pattern in the parameters is explainable. These results should be validated within the framework of the three classes of tests described previously.

To remove this constraint in the linear predictor, it is common to introduce interaction terms. An interaction term is a non-linear construct that allows the results of one predictor to become dependent on the value of another predictor. Building on the prior example, assume we have the following model structure:

$$\text{Base} + \text{RatingArea} + \text{NamedInsuredAge} + f_1(\text{RateChange}) + f_2(\text{CompetitivePosition}) + f_2(\text{CompetitivePosition.RatingArea})$$

For this model structure the curve parameters for the competitive position will vary based on rating area. Specifically:

Level	Predictor	Component
Base		-0.37781
NamedInsuredAge	25	0.38846
RatingArea	1A	0.09457
RateChange	1.09	-0.9216 * 1.09 + 1.0368
CompetitivePosition	1.07	-0.9234 * 1.07 + 1.1542
CompPosit.RatingArea	1A - 1.07	-0.4548 * 1.07 + 0.5685
Linear Predictor		0.38550
Fitted Value		0.59520

Table 2.2 Model structure example; interactions between competitive position and rating area

Similar to the discussion related to the main effect, adding interaction terms to the model should be done within the context of model testing. Identifying interactions are a great way to complicate the model structure to improve its explanatory power. This process is generally a non-trivial task. The various strategies for identifying interactions are outside the scope of this paper.

Recall the main purpose in building a demand model is to project the expected price sensitivity of the in-force customer or new business quote – in other words, the elasticity. Using the definitions from the prior section, we see that the elasticity is a slope function associated with the price factors in the model. Oftentimes, when building a model structure, we are looking for interactions (i.e. non-linear relationships) between the price and non-price factors. Within a traditional linear predictor, the analyst could build a model structure where all price factors interact with all non-price factors. This model would have a strong chance of over-fitting the data, and there is a possibility that a derived negative elasticity could occur from such a model.

Arbitrarily adding interactions is generally not a good idea. This could lead to over-fitting which would weaken the predictive power of the model. In addition, adding any type of complexity to a model should be validated based on statistical, consistency and judgment tests. As in other models, if a structure does not produce coherent results, then that structure should not be included in the model form. However, there are technical solutions available to build these complex interactions forms without resulting in the negative elasticity issue.

Localization

The first approach is to build localized models. Essentially, the analyst would use a decision tree tool to split the data into various subsets. The local main effects models will be built for each subset. By splitting the data into more homogenous subgroups, the modeler can build simpler models on smaller sets of data which would have the same effect of building complex models on larger sets of data. The localization procedure is easy to explain and implement; however, depending on the split, the model could produce discontinuities in some of the continuous factors. For example, if the split was based on rate increases versus rate decreases, then the resulting elasticity projections around the no rate change would not likely be continuous.

Customer Scoring

The second approach is to use a form of customer scoring to build the model structure. In this approach, a main effect model is built to include price and non-price factors. Interactions among the non-price factors are then incorporated as necessary. The parameters from the non-price factors are then combined into a customer score. This score is then interacted with the price factors to yield the necessary complexity. Clarifying this with an example, assume we have the following model form:

$$\text{Base} + \text{RatingArea} + \text{NamedInsuredAge} + f_1(\text{RateChange}) + f_2(\text{CompetitivePosition})$$

Using the parameters from RatingArea and NamedInsuredAge, the following model form is built:

$$\text{Base} + \text{Score} + f_1(\text{RateChange}) + f_2(\text{CompetitivePosition})$$

Where Score = RatingArea + NamedInsuredAge

Note at this point both of these models would produce the exact same fitted values. The final step is to introduce the interaction:

$$\text{Base} + \text{Score} + f_1(\text{RateChange}) + f_2(\text{CompetitivePosition}) + \text{Score}.f_1(\text{RateChange}) + \text{Score}.f_2(\text{CompetitivePosition})$$

The following picture illustrates this effect:

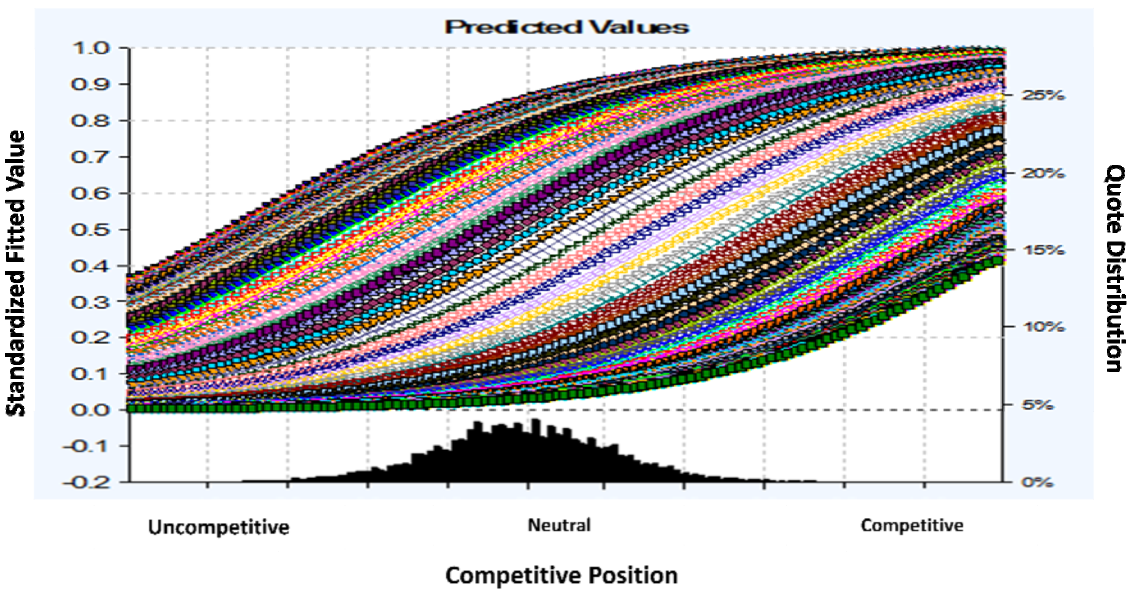

Figure 2.6 Interaction between customer score and competitive position

The x-axis represents the competitive position defined as MarketPrice/CompanyPrice. The MarketPrice is defined as a complex weighted average of a collection of competitors' prices. The left hand side of the x-axis represents less competitive rates and the right hand side represents more competitive rates. The left hand y-axis is the fitted value from the model and the right hand y-axis is the quote distribution. The score of the customer is derived from the non-price factors as described earlier. Sometimes this is labeled as the loyalty of a customer because it reflects the component of the expected demand that is not based on price. Each line represents a customer score where the higher the score the greater the 'loyalty' of the customer. When this customer score is interacted with the competitive price factor, we see that the steepest lines (i.e. most price sensitive) are evident in the middle of the tapestry. This makes sense because the most loyal customers (the flatter lines on the top end of the tapestry) are less price sensitive as they identify with either the brand or other services provided by the company. On the opposite end, the least loyal customers (the flatter lines on the bottom end of the tapestry) are less price sensitive as they do not identify with the company brand or do not seem to be interested in the additional services provided by the company. The middle of the spectrum represents customers who are more influenced by the pricing and less so by branding.

The customer scoring approach allows the modeler to introduce a significant level of complexity into the model form. There are interpretative challenges of what makes a more loyal customer for the non-price factors. In addition, the two staged approach implies an amount of quantitative certainty associated with parameters from the original main effects model.

Non-linear Modeling

The third approach defines the scoring parameters and the main effects parameters within the same fitting algorithm. The following nonlinear model form is constructed:

$$\text{Base} + \text{RatingArea} + \text{NamedInsuredAge} + f_1(\text{RateChange}) + f_2(\text{CompetitivePosition}) + \text{Score}.f_1(\text{RateChange}) + \text{Score}.f_2(\text{CompetitivePosition})$$

Where Score = exp(RatingArea + NamedInsuredAge)

Essentially, there are two sets of non-price coefficients that are being calculated within the model. The first represent the parameters to be applied to the main effects. The second set correspond to the embedded score function. This solution is different from the customer scoring approach because it is based on deriving the score parameters at the same time as the main effect parameters. Recall the customer scoring solution is a recursive process where the customer score is parameterized in the main effect. This score is then introduced in a subsequent model both as a main effect and an interaction term. The advantage of the non-linear model is to remove the element of recursion. Note that the more stable the main effects parameters, the less the difference between the prior two model forms.

As with the customer scoring approach, there are interpretive challenges to the non-linear structure. In addition, the specific structure can produce convergence issues in the fitting algorithm; however, this can be partly alleviated by simplifying the structure. Finally, this approach (along with the customer scoring strategy) ignores the potential for real negative elasticity. This possibility can exist when a company takes massive rate decreases resulting in cancellations. When customers see a large decrease, they tend to believe that they were previously and unfairly priced too high. This jeopardizes the customer/insurer relationship and causes many customers to shop and switch insurers.

2.4. Validation

The resulting models from these approaches need to be validated using out of time samples to test their performance. In earlier sections, validation techniques which were associated with specifying the correct link function were discussed. In this section, we will focus on two additional validation strategies associated with a validation data set.

The validation data set can either be a random hold out sample from the original data OR an out of time sample OR both. The value of an out of time sample is that it reflects an accurate measure of predictive accuracy. The models built on historical data are applied to subsequent data. What better way to assess the model performance than to use 'future' data. The disadvantage of an out of time sample is that there could be a significant change in the way business operates that is not reflective of the historical data (e.g. sudden increased presence of internet aggregators). So, the lack of predictive power is not a fault of the model, but rather due to the fact that the times are changing. A similar converse statement can be made about a random sample – it is a good measure of explanatory power and will not be skewed by distributional drifts in the trend; however, the measure provides a weak insight into how well the model will perform in the future. Ultimately, the modeler decides on variations of these two themes.

Once the validation data set has been defined, the model structure and resulting parameters that were built from the training and testing data can be used to score the validation records. These scores can then be compared to the actual values to assess model performance. Consider the following example:

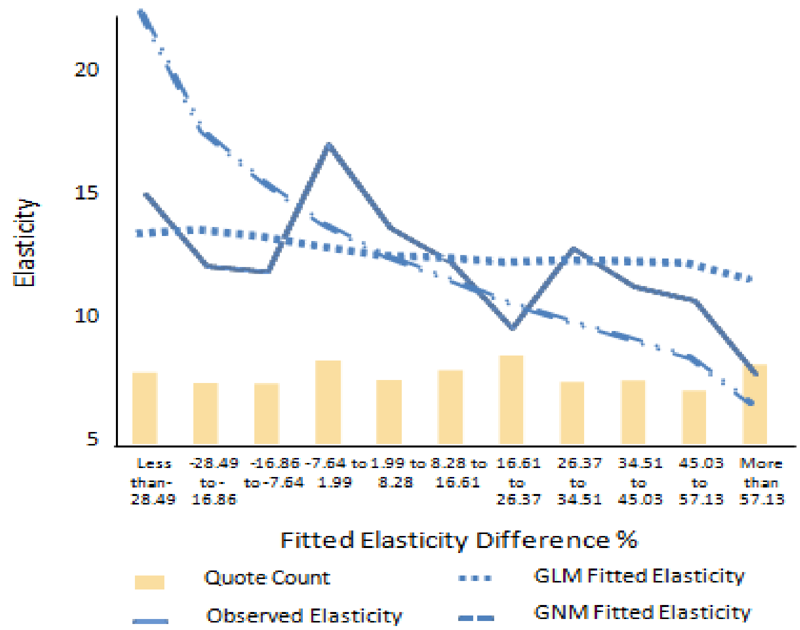

Figure 2.7 Validation of elasticity models against observed elasticity

In this example, two models were compared to actual results based on an out of time validation data set in much the same way that was described in the previous section. As before, the out of time data is the number of quotes from months seven and eight. The first model (dotted) was built using a standard logistic regression within a GLM framework. The model form was localized and several interactions were introduced into the localized components. The second model (dashed) was built using the non-linear method described in the prior section. Applying the structure and parameters to the out of time validation data set, a modeled elasticity score was generated from both models. The quote data was then ranked into approximately 11 equal weighted buckets based on the ratio of the score from both models. Specifically:

(GLMFittedElasticty/GNMFittedElasticity – 1) * 100

This ratio represents the x-axis in Figure 2.7. For each of these buckets, the aggregate modeled elasticities were calculated as described earlier (the change in close rate divided by the change in price). The dashed line and the dotted line intersect when both models produce the same score. To the left of the intersection, the dashed line is higher than the dotted. This represents the set of observations in which the GNM produces a higher expected elasticity. To the right of the intersection, the dashed line is less than the dotted which represents when the GNM approach produces a lower expected elasticity.

Finally, the observed elasticities (solid line) for each of these buckets were plotted. Comparing these different lines allow the modeler to assess which structural form is doing a better job of predicting the elasticity. In this particular case, it appears as if the non-linear model structure is over-fitting the historical data. Thus, it has weaker predictive power. This result is more apparent when the GNM predicts a higher elasticity than the GLM.

In the prior example, the model structure and resulting parameters were validated on an out of time sample. From this approach, we can make general statements about the appropriateness of the model form; however, we are unable to make specific statements about the strengths and weaknesses of particular aspects of the model structure. For example, how predictive is the set of parameters coming from a particular interaction? An approach to deal with this problem is to use the validation data to reparameterize the model structure that was built from the historical data. The modeler can then compare the parameters calculated from the training and testing data to the parameters calculated from the validation data. The inconsistencies between the sets of the parameters indicate where there are weaknesses in the predictive power of certain aspects of the model structure.

Modeling the demand and deriving elasticity provides a powerful understanding of the customer's reaction to the presented price. This understanding is crucial in setting the right rate for the risk. These models require different technical expertise than those of standard costs models, as was described by their structure and form. Ultimately, when building a model it is imperative to realize that, regardless of strategy, the best models tend to be a mixture of technical expertise combined with business acumen.

3. PRICE SIMULATION

In the following two sections, we will look at practical applications of PED, beginning with price simulation. With the help of our demand models, price simulation allows us to "predict" how our book of business will change over time with respect to key metrics such as profitability and demand. To gain an understanding of how price simulation can be utilized in the pricing process, the remainder of this section will focus on a specific example.

3.1. Scenario Testing

During a rate review, it is not uncommon for an insurance company to have a few different scenarios that it would like to test. Although the company may have a high-level idea in mind (such as a 5% rate decrease), there are a number of ways that 5% decrease can be allocated. Price simulation allows us to determine which scenario will help us best achieve our goals. For example, suppose an insurer writes personal auto insurance. Based on initial countrywide planning efforts, the insurer has decided to pursue a 5% rate decrease in state X. In addition, it would like to simulate two price scenarios to help determine which one should be implemented. The scenarios are as follows:

- Scenario 1 – 5% base rate decrease
- Scenario 2 – 15% decrease for operators aged 25-30 off-balanced to an overall 5% decrease

Once the scenarios have been determined, the assumptions that feed the simulator must be developed. For this example, assume the following:

- Datasets
  - Quotes
  - Renewals
- Conversion Model
  - Price Component - Premium
- Retention model
  - Price Component – Market Competitiveness (fixed)
- Time Horizon – Four periods, each lasting six months
- Quote Growth Rate – 5% each period
- Quotes do not enter simulation until new rates go into effect at the beginning of period 1
- Quote distribution constant over time (i.e. if 21 year olds represent 15% of the quotes in period 0, they will represent 15% in subsequent periods as well)
- Aging Assumptions
  - Operators age by 1 every other period
  - Vehicles age by 1 every other period

3.2. Running the Simulation

Now that the assumptions have been identified, the insurer can run a simulation. For educational purposes, we will examine "quotes" and "renewals" separately to gain an understanding of the mechanics underlying price simulation. Let us begin with "quotes":

Quotes								
	(1) Period	(2) Policies Offered	(3) Policies Written	(4) = (3) / (2) Conversion Rate	(5) Policies Retained	(6) = (5) / (3) Retention Rate	(7) Profit Margin	(8) Elasticity
Scenario 1	0	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	1	20,000	5,493	27.5%	4,669	85.0%	1.9%	1.8
	2	21,000	5,767	27.5%	4,902	85.0%	1.9%	1.8
	3	22,050	6,058	27.5%	5,150	85.0%	1.9%	1.8
	4	23,153	6,360	27.5%	5,406	85.0%	1.9%	1.8
Scenario 2	0	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	1	20,000	5,646	28.2%	4,743	84.0%	1.8%	2.4
	2	21,000	5,928	28.2%	4,980	84.0%	1.8%	2.4
	3	22,050	6,228	28.2%	5,231	84.0%	1.8%	2.4
	4	23,153	6,538	28.2%	5,492	84.0%	1.8%	2.4

Table 3.1 Quote simulation

In period 0, there are no **Policies Offered** since we assumed quotes do not enter the simulation until period 1. In period 1, the initial quote dataset of 20,000 enters the simulation. In period 2, 21,000 new quotes enter the simulation based on the assumed 5% quote growth rate. Periods 3 & 4 follow similarly. **Polices Written** shows the number of quotes converted based on the assumed conversion model. Once policies written are determined, an aggregate **Conversion Rate** can be calculated, as shown in Column (4). **Policies Retained** gives the number of converted quotes that persist to the next period based on the assumed retention model. Once policies retained are determined, an aggregate **Retention Rate** can be calculated, as shown in Column (6). Column (7) shows the **Profit Margin** changes over time in each scenario. The conversion rate, retention rate and profit margin are the same for each period because the quote distribution remains constant over time per our assumptions. Lastly, Column (8) contains the aggregate elasticity for the quote dataset. Since elasticity is a function of demand and demand is higher in Scenario 2, it follows that the Scenario 2 elasticity is higher than the Scenario 1 elasticity. Based on Table

3.1, which scenario is better? From a growth perspective, Scenario 2 is superior. However, from a profit margin perspective, the insurer would recommend Scenario 1. Before making a final decision, we must consider "existing business":

Renewals					
	(1) Period	**(2)** Policies Offered	**(3)** Policies Retained	**(4) = (3) / (2)** Retention Rate	**(5)** Profit Margin
Scenario 1	0	50,000	44,000	88.0%	2.5%
	1	44,000	41,287	93.8%	2.4%
	2	45,956[1]	44,162	96.1%	2.3%
	3	49,064	47,147	96.1%	2.2%
	4	52,296	49,315	94.3%	2.2%
Scenario 2	0	50,000	44,000	88.0%	2.5%
	1	44,000	41,287	93.8%	2.4%
	2	46,030	44,155	95.9%	2.5%
	3	49,135	47,121	95.9%	2.6%
	4	52,352	49,263	94.1%	2.7%

Table 3.2 Renewal simulation

In period 0, policies offered equal the initial renewal dataset of 50,000 policies. In period 1, policies offered equal the policies retained from period 0. This number is the same in both scenarios since we assumed the price component in the retention model was fixed (i.e. the market price and competitor price were perfectly correlated). In period 2, converted quotes from period 1 have transitioned to renewals. Thus, policies offered equal the sum of period 1 policies retained from the quotes in Table 3.1 and the renewals table in Table 3.2. Periods 3 and 4 follow similar logic. **Policies Retained** shows the number of renewal policies that are retained to the next period based on the assumed retention model. Once policies retained are determined, an aggregate **Retention Rate** can be calculated, as shown in Column (4). Lastly, Column (5) shows how the **Profit Margin** changes over time in each scenario. Based on Table 3.2, which scenario is better? From a retention perspective, Scenario 1 is best. However, from a profit margin perspective, the insurer would choose Scenario 2. Since Tables 3.1 and 3.2 failed to provide a consistent solution, we must examine the total results to determine which scenario to recommend.

[1] 45,956 = 41,287 + 4,669

Total (Renewals + Quotes)	(1) Period	(2) Policies Offered	(3) Policies Written	(4) Policies Retained	(5) Earned Premium	(6) Profit Margin	(7) Absolute Profit
Scenario 1	0	50,000	50,000	44,000	$35,250,000	2.5%	$881,250
	1	64,000	49,493	45,956	$34,486,258	2.3%	$810,152
	2	66,956	51,723	49,064	$36,412,258	2.3%	$822,930
	3	71,114	55,122	52,296	$38,800,399	2.2%	$842,146
	4	75,449	58,657	54,722	$40,949,798	2.2%	$888,759
Scenario 2	0	50,000	50,000	44,000	$35,250,000	2.5%	$881,250
	1	64,000	49,646	46,030	$34,729,064	2.3%	$812,026
	2	67,030	51,958	49,135	$36,692,114	2.4%	$891,271
	3	71,185	55,363	52,352	$39,087,466	2.5%	$985,029
	4	75,505	58,890	54,755	$41,236,423	2.6%	$1,076,159

Table 3.3 Total business simulation

After viewing the combined results in Table 3.3, the answer is clear. Over the stated time horizon of two years, Scenario 2 outperforms Scenario 1 from growth, retention, and profit perspectives. Whereas traditional analysis provides little distinction between the scenarios, price simulation allowed the insurer to "predict" how its book will change over time in each scenario, resulting in the better choice.

Please note that the above example is intended to illustrate the concept of price simulation. With a slight difference in rate allocation, we see that Scenario 2 yields more policies (~30) and higher profit (~$400,000) than Scenario 1 by the end of period 4. Not surprisingly, further rate allocation with a larger book of business could lead to a more significant difference in pricing scenarios.

4. PRICE OPTIMIZATION

Suppose ABC Insurance wants to increase its new business conversion rate. One option is to decrease the overall premium level. This will certainly increase the company's conversion rate, but it's probably not the optimal solution because it affects each insured equally. Another approach is for the company to be more surgical in its approach by targeting specific market segments with higher growth potential. Although this option is superior to the first, current techniques rely on sub-optimal human judgment when deciding which segments to target. Price optimization is a mathematical procedure that allows insurers to determine the appropriate rates to charge to maximize some metric over some specified time horizon. As stated above, ABC Insurance wanted to increase its overall conversion rate on new business. Although the approaches described will do that, the result will be inferior to a price optimization routine that analyzes thousands of different pricing options. Examples of metrics a company might want to maximize include:

- Profitability
    - Percentage basis (margin)
    - Absolute basis (total dollars)
- New Business Conversion
- Renewal Business Retention
- Competitive Position

It is important to note that these metrics are not mutually exclusive. Typically, a company sets a constraint on one metric while maximizing another. For example, a company may want to maximize its new business conversion rate while maintaining its current profitability. Similarly, a company may want to maximize its profitability while maintaining its current conversion rate.

4.1. Price Optimization Methods

In this section, we will discuss two price optimization methods. The first method optimizes prices at the rating structure level by determining optimal rating factors for the structures the company has chosen to revise. This method optimizes on the current rating structure directly, and is therefore easy to implement. However, it fails to identify gaps in the current rating structure where growth potential is not being realized. The second method optimizes prices irrespective of the current rating structure by determining an optimized premium for each insured in a firm's book of business. This has a number of advantages:

- Optimizes premium at the individual insured level
- Provides opportunity for improvements in a firm's rating structure by optimizing rates outside of the current pricing system
- Produces an efficient frontier of optimized premiums representing the highest possible demand at varying profit levels

Despite these advantages, this method has a few drawbacks. First, it requires more time since a rating structure must be reverse engineered from the optimized premiums. Second, some of the benefit gained from the optimization routine is lost during the reverse engineering process due to modeling limitations. In addition to the disadvantages listed above, both methods suffer from regulatory constraints, which will be discussed in a later section. The remainder of this section will focus on the second method.

4.2. The Benefit Function

Assume a firm wants to maximize its new business conversion rate while maintaining its current profitability. The firm has decided to use price optimization to achieve this goal. In order to optimize prices at the individual insured level, we begin with the benefit function. In basic terms, the benefit function calculates the expected profit for each insured over a specific time horizon. In equation form, the one-year benefit function might look like this:

$$BF_i = CD_i * (Q_i - L_i - E_i) \tag{4.1}$$

Where BF = Benefit Function

CD = Cumulative Demand

Q = Proposed Premium

L = Pure Premium

E = Expenses

i = i^{th} insured

Let us examine each piece of the equation starting with cumulative demand. In order to calculate **cumulative demand**, we need a conversion model and a retention model. Each of these models can be created using the techniques described in Section 2 of this paper. Once the models are built, we simply accumulate total demand for each individual over the stated time horizon. For a one-year time horizon, the cumulative demand for new business is the probability that the firm acquires the insured in year one. If the time horizon were three years, then the cumulative demand for new business would be the probability that the firm acquires the insured in year one, multiplied by the probability that the firm retains the insured for two more years. The **proposed premium** is derived by solving for the optimized premium. The **pure premium** is an estimate of the expected loss costs for each insured. Lastly, we must account for **expenses**. Once we have defined each piece of the benefit function, we can solve for the proposed premium (i.e. optimized premium) by maximizing the benefit function. This results in an optimized premium for each insured that maximizes his or her expected profit over the stated time horizon irrespective of the conversion rate. Although this certainly qualifies as an optimized premium, it is not the optimized premium sought by the firm in our example. Remember, our firm wanted to increase its new business conversion rate **while maintaining its current profitability.** So, where does the firm go from here?

4.3. The Efficient Frontier

The efficient frontier is a curve that plots the expected loss ratio against demand. It is "efficient" because each point on the curve represents a set of optimized premiums. The maximized benefit function exists on the efficient frontier at the point where profitability is the greatest. As we saw earlier, it has a set of "optimized" premiums associated with it; namely, the ones that maximize profitability over the stated time horizon. As you move up the curve, demand increases and profitability decreases. Points inside the curve, such as the firm's current position, are considered sub-optimal because the company could improve demand while maintaining profitability, improve profitability while maintaining demand, or some combination thereof. If a firm wants to maximize its new business conversion rate while keeping its current profitability constant, the firm moves its position up the curve until it reaches the optimized location. To move up the curve, the optimization routine analyzes the PED curve for each individual insured to determine the proper price to charge them in order to achieve the target expected loss ratio.

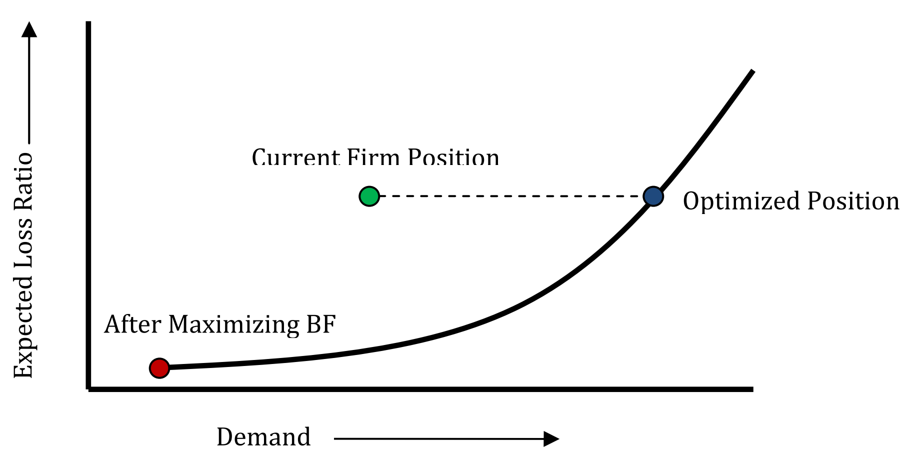

Figure 4.1 The Efficient Frontier

Once a set of optimized premiums has been determined that meets the firm's goals, the firm can finally reverse engineer a rating structure by modeling the optimized premiums. In general, this just involves fitting the current rating structure onto the optimized premiums. However, thanks to the flexible nature of optimizing at the individual insured level, other predictors may be used to account for gaps in the current rating structure found during the reverse engineering process. Although there are a number of ways to complete the reverse engineering process, we have chosen to omit that discussion due to the limited scope of this paper.

4.4. Implementing Optimized Rates in a Regulatory Environment

One substantial challenge associated with price optimization is implementing optimized rates within the constraints of a regulatory environment. The important thing to remember is that price optimization is a tool designed to help firms make better pricing decisions. As with any rate filing, a company starts by developing indications based on the loss propensities, expense loads and rate of return objectives for various segments of its book of business. At this point, the company can choose to stay at the current rates, use the indicated rates, or deviate from the indicated rates based on business judgment. In this regard, price optimization can be seen as a deviation from the indicated rates. Of course, this doesn't mean that a firm can file anything that comes out of the optimization routine. Since price optimization relies on PED in addition to loss propensity, it could conflict with traditional pricing methods. For example, suppose a mature segment of a firm's book of business has a traditional actuarial rate indication of -10%. In addition to this, assume the segment has a low PED relative to the rest of the book. Now, assume the firm runs an optimization routine designed to increase overall demand. Due to the low PED of the mature segment, the optimized rate will most likely be higher than the current rate since that particular segment is less price sensitive. However, this conflicts with the indicated decrease of -10%. In this case, the optimized rates could be subject to regulatory challenge and may need to be revised to coincide with the loss indications. As a result, some of the benefit from the optimization routine would be lost. This example is not meant to deter the use of price optimization; it's meant as a reminder that the final rates decided upon still require actuarial justification and approval by state agencies.

5. CONCLUDING REMARKS

In this paper, we detailed basic techniques for developing demand models with emphasis on model form and model structure. Next, we examined the use of PED in the context of two practical applications: price simulation and price optimization. Lastly, we described how to implement optimized rates in a regulatory environment.

Despite the sophistication present in the insurance industry today, the integration of competitive data within the pricing process continues to invite further exploration. This paper provides the reader with the basic tools needed to begin this process. With these tools, insurers possess the necessary knowledge and skills to implement more targeted pricing decisions that better achieve their goals.

Acknowledgments

We would like to thank the following individuals for their comments and suggestions: William Frierson, Rob Kane, David Maurer, & Ryan Voge. We would also like to give special thanks to Rob Kane for his numerous contributions to the price simulation section and to Michael Korner for his invaluable and ongoing research in demand modeling techniques.

6. REFERENCES

1. Anderson, Duncan, et al. "A Practitioner's Guide to Generalized Linear Models." *CAS Discussion Paper Program* (2004): 1-115.

2. Brockman, Michael J. "Statistical Motor Rating: Making Effective Use of Your Data." *Journal of the Institute of Actuaries* (April 1992): 457-543.

3. Hastie, T., R. Tibshirani, and J.H. Friedman. *The Elements of Statistical Learning*. Springer, 2001.

4. Holler, Keith D., David Sommer, and Geoff Trahair. "Something Old, Something New in Classification Ratemaking with a Novel Use of GLMs for Credit Insurance." *CAS Forum* (Winter 1999): 31-84.

5. McCullagh, P., and J.A. Nelder. *Generalized Linear Models*. 2nd ed. Chapman and Hall, 1989.

6. Mildenhall, Steve. "Systematic Relationship Between Minimum Bias and Generalized Linear Models." *CAS Proceeding LXXXVI* (1999): 393-487.

7. Murphy, Karl P., Michael J. Brockman, and Peter K. Lee. "Using Generalized Linear Models to Build Dynamic Pricing Systems." *CAS Forum* (Winter 2000): 107-139.

Author Biographies

- **Serhat Guven** is a Senior Consultant with Towers Watson. He is a Fellow of the CAS and a member of the American Academy of Actuaries. He can be reached at serhat.guven@towerswatson.com.
- **Michael McPhail** is a Pricing Director with USAA. He is a Fellow of the CAS and a member of the American Academy of Actuaries. He can be reached at michael.mcphail@usaa.com.